CEWES MSRC/PET TR/98-50

User Impact Study: CEWES MSRC CRAY C90 Decommissioning

by

Henry A. Gabb

DoD HPC Modernization Program
Programming Environment and Training
CEWES MSRC

CEWES Major Shared Resource Center MSRC

Nichols Research

11h00198

Work funded by the DoD High Performance Computing Modernization Program CEWES Major Shared Resource Center through

Programming Environment and Training (PET)

Supported by Contract Number: DAHC 94-96-C0002
Nichols Research Corporation

User Impact Study: CEWES MSRC CRAY C90 Decommissioning

Henry A. Gabb
CEWES MSRC, CMG Lead

Introduction

The CEWES MSRC CRAY C90 (C90) is scheduled for decommissioning sometime in 1999. It is no longer a viable high performance computing (HPC) platform at this site due to its high maintenance costs and poor price-to-performance ratio. Many problems that were HPC applications five years ago will run quite well on a high-end workstation now. Today, large HPC problems are better served by modern scalable parallel architectures.

However, prejudices against parallel computing exist within the scientific community. Many scientists and engineers believe that parallel programming is too difficult or does not merit the cost in manpower. There is also the belief that parallel computing in the last decade has failed to deliver the performance promised. This perception, though true at one time, is now false.

Writing parallel programs is more complex than writing serial programs, but standardization of parallel programming models (e.g., message passing and threads)[1-3] has greatly improved both usability and portability. Two new data parallel Fortran dialects, High Performance Fortran[4] and Co-array Fortran (formerly known as F--),[5,6] as well as automatic parallelizing compilers from some vendors, are also available. The high performance computing industry has also adopted a unified set of parallel compiler directives to facilitate parallel programming on shared memory architectures.[7-9] These recent advances combine to make parallel computers significantly more manageable to the typical scientific programmer.

A migration effort is already underway at the CEWES MSRC. The Computational Migration Group (CMG) assists C90 users in migrating their codes to the scalable parallel systems. CMG assistance takes many forms: from providing technical documentation to direct involvement in code porting and parallelization.[10] Technical documentation includes the Frequently Asked Questions documents and technical reports produced by the CMG.[11-20] The technical reports cover issues in performance, parallelism, and porting C90 codes.[11-20] The CMG also maintains a Fortran module of functional equivalents to many Cray vector intrinsics. For many users, however, the CMG is simply a base of scientific, mathematical, and general computing expertise.

The last year of operation for the C90 is 1998, and one question remains: Is the current migration effort sufficient? An impact survey was sent to all C90 account holders on Friday, 28 August 1998. This survey was designed to determine the readiness of the CEWES MSRC user community to migrate to scalable parallel systems. Seventy-four users responded. The results are analyzed in this report, and the actual survey is included as an appendix.

Resistance to Migration?

Nearly one-third of those responding to this survey are already working on a CEWES MSRC parallel computer. Most of those not currently working on a parallel platform express an interest in migrating. It appears that users are willing to move to a parallel machine when the C90 is removed.

A surprising number (10 percent) are unsure whether to migrate to a parallel computer. There are several possible explanations for this uncertainty: skepticism that parallelism will be beneficial, or an unwillingness to devote time and resources to parallelizing working programs. The CMG and the

Programming Environment and Training Program are both responsible for educating these users about the advantages of parallel computing.

Commercial Software

Only about 20 percent of respondents are currently using commercial software. A list of commercial applications currently being used on the C90 was obtained from this survey. Of the commercial products listed, ABAQUS, NASTRAN, CTH, ENSIGHT, and TRUEGRID already have versions available on at least one CEWES MSRC parallel system. GASP is another heavily used application for which a parallel version exists (for the SGI Origin 2000), but it is not available at CEWES MRSC. As the decommissioning date approaches, the parallel version of GASP should be obtained.

One key application, DISSPLA, does not have a parallel version. An effort is underway to find a suitable replacement. In cases where an important piece of commercial software is not available for at least one of the CEWES MSRC parallel systems, the CMG will have to take steps to ensure that the affected users are able to continue working when the C90 is removed. This could mean finding a Cray vector system at another site or finding functionally equivalent software. Within the Department of Defense High Performance Computing Modernization Program (HPCMP), the Common High Performance Computing Software Support Initiative (CHSSI) exists to develop and maintain scalable parallel versions of forty key applications across ten computational technology areas. CHSSI codes can probably replace some commercial applications.

Custom Developed Software

The majority of users are running custom developed applications. Over half of the respondents (60 percent) use custom software on the C90. The majority of these respondents indicate that their codes take advantage of the Cray vector architecture. Over three-quarters say that their codes are vectorized, and approximately half utilize the Cray vector intrinsics. It is safe to assume that most commercial applications are also vectorized to some degree. It should be noted, however, that questions 5a and 5b do not address the degree of vectorization. The PerfStat performance monitor (Instrumental, Inc.) is being used to precisely determine vector utilization on the C90.

Reliance on vector processing does not pose a significant barrier to migration. For example, most Cray vector intrinsics are available on the SGI Origin 2000 or CRAY T3E. Some are even available on the IBM SP. The CMG also maintains a Fortran 90 module containing functional equivalents to many Cray vector intrinsics. However, performance loss is an important consideration for extensively vectorized codes. Codes that are not vectorized will probably show similar performance on the scalable parallel computers. On the SGI Origin 2000, automatic parallelization is likely to give improved performance with very little effort.

The typical custom applications running on the C90 are not memory bound (question 5c). A two-gigabyte threshold was chosen because this is the upper memory limit on the CEWES MSRC IBM SP nodes. This also tends to be the upper memory limit for workstations. Therefore, a workstation or single IBM SP node migration path is a viable option for many applications currently running on the C90. The single-node IBM SP option gives users a place to run their applications at the CEWES MSRC while the source code is being parallelized. Most users responded that their applications would benefit from parallel processing.

Outlook for the Computational Migration Effort

The results of this survey are encouraging to the CMG. Most C90 users are either working on a parallel machine or express an interest in moving to a parallel environment. Of the commercial applications identified by users, many already have parallel versions (ENSIGHT, ABAQUS, TRUEGRID, DYNA2D, GASP, CTH) or have parallel versions pending (NASTRAN, DYNA3D). Of the users running custom developed applications, most take advantage of the vector architecture. Whereas use of vector intrinsics does not pose a problem, performance loss in vectorized codes is still a concern. The overall conclusion drawn from this survey is that C90 users are willing to move to parallel environments and in many cases have done so on their own initiative. Therefore, a successful migration effort in 1999 is not an insurmountable task.

References

1. Snir, M., Huss-Lederman, S., Walker, S., and Otto, S. W. (1996). *MPI: The Complete Reference*. The MIT Press.

2. Geist, A., Beguelin, A., and Dongarra, J. (1994). *PVM: Parallel Virtual Machine*. The MIT Press.

3. Nichols, B., Buttlar, D., and Farrell, J. P. (1996). *Pthreads Programming*. O'Reilly and Associates, Inc.

4. Koelbel, C. H., Loveman, D. B., Schreiber, R. S., Steele, G. L., and Zosel, M. E. (1994). *The High Performance Fortran Handbook.* The MIT Press.

5. Numrich, R. W. and Steidel, J. L. (1997). "F--: A simple parallel extension to Fortran 90," *SIAM News.* 30(7),1-8.

6. Numrich, R. W. and Reid, J. K. (1998). "Co-array Fortran for parallel programming," *ACM Fortran Forum*.

7. "OpenMP: A proposed industry standard API for shared memory programming," OpenMP Architecture Review Board. www.openmp.org (October 1997).

8. "OpenMP Fortran application program interface (Version 1.0)," OpenMP Architecture Review Board, www.openmp.org (October 1997).

9. Clark, D. (1998). "OpenMP: A parallel standard for the masses," *IEEE Concurrency.* 6(1),10-12.

10. Bova, S. W., Gabb, H. A., and Stagg, A. K. (1998). "Practical aspects of migrating DoD codes to scalable architectures," CEWES MSRC/PET TR/98-37.

11. Nagle, D. (submitted). "Real v. double precision on various machines," CEWES MSRC/PET TR.

12. Nagle, D., (submitted). "Using Fortran procedure interfaces," CEWES MSRC/PET TR.

13. Nagle, D., (submitted). "Floating point arithmetic," CEWES MSRC/PET TR.

14. Nagle, D., (submitted) "Cray pointers v. Fortran 90 pointers," CEWES MSRC/PET TR.

15. Nagle, D., (submitted). "Fortran 95 pure procedures and elemental procedures," CEWES MSRC/PET TR.

16. Nagle, D., (submitted). "Porting the Fortran module standard_types," CEWES MSRC/PET TR.

17. Nagle, D., (submitted). "Using the Fortran module standard_types," CEWES MSRC/PET TR.

18. Nagle, D., (submitted). "Fixed format to free format conversion," CEWES MSRC/PET TR.

19. Nagle, D., (submitted). "Fortran 77 v. Fortran 90 namelist usage," CEWES MSRC/PET TR.

20. Nagle, D., (submitted). "Cache blocking and derived types on the CEWES MSRC parallel computers," CEWES MSRC/PET TR.

21. Bova, S. W., Breshears, C. P., and Gabb, H. A. (1998). "Status report on parallelization of MAGI," CEWES MSRC/PET TR/98-49.

Appendix A
Survey and raw data

Subject: User Requirements Survey for the CRAY C90

As a CEWES MSRC user you understand the requirement to have the right hardware and software available to do your work. With this in mind the CEWES MSRC is gradually shifting its high performance computing emphasis towards scalable parallel supercomputers. Many of the large computational problems encountered in modern scientific research are only feasible using parallel computation. The Computational Migration Group was formed at CEWES MSRC to help users migrate their C90 codes (or provide functionally equivalent parallel applications) to the scalable parallel computers. To determine the future requirements of our CRAY C90 user community, we ask that you complete the following survey.

Name:
Project ID:
Project title:
Principal investigator:

Question	Yes	No	Don't Know
1. Are you currently working on any of the CEWES MSRC parallel computers?	23	51	0
2. Are you interested in migrating to a parallel computer?	50	17	7
3. Do you use any commercial software on the CRAY C90? (If yes, please list.)*	16	56	2
4. Do you run custom developed applications on the CRAY C90?	44	30	0
5. If you run custom developed applications on the CRAY C90, please answer the following four questions			
a. Are these applications vectorized?	34	9	1
b. Do these applications use the Cray vector intrinsics?	21	18	5
c. Do these applications require more than two Gbytes of memory?	8	32	4
d. Would these applications benefit from parallel processing?	29	6	9

*ABAQUS (5), DISSPLA (3), GASP (3), ENSIGHT (2), NASTRAN, CTH, DYNA2D, DYNA3D, MINT, ACES2, HISTORIAN, ALE3D, TRUEGRID